

Reconfiguring Ordered Bases of a Matroid

Anna Lubiw*

Vinayak Pathak*

December 6, 2016

Abstract

For a matroid with an ordered (or “labelled”) basis, a basis exchange step removes one element with label l and replaces it by a new element that results in a new basis, and with the new element assigned label l . We prove that one labelled basis can be reconfigured to another if and only if for every label, the initial and final elements with that label lie in the same connected component of the matroid. Furthermore, we prove that when the reconfiguration is possible, the number of basis exchange steps required is $O(r^{1.5})$ for a rank r matroid. For a graphic matroid we improve the bound to $O(r \log r)$.

1 Introduction

Broadly speaking, “reconfiguration” is about changing one structure to a related one via discrete steps. Examples include changing one triangulation of a point set in the plane to another via “flips”, or changing one independent set of a graph to another by adding/deleting single vertices. More broadly, other examples are to sort a list by swapping pairs of adjacent elements, or to change a basic solution of a linear program to an optimum one via pivots. Central questions are: to test whether one configuration can be changed to another using the discrete steps; and to compute, or bound, the number of steps required. Reconfiguration problems have been considered for a long time, and explicit attention was drawn to complexity issues by Ito et al. [5] and by van den Heuvel [4].

A fundamental result about matroids is that any basis B can be reconfigured to any other basis B' using a sequence of *basis exchange* steps. In each *basis exchange* one element of the current basis is replaced by a new element to yield a new basis. The number of basis exchange steps needed to reconfigure B to B' is the difference $|B - B'|$.

In this paper we explore a “labelled” version of basis reconfiguration where the elements of each of the two initial bases are labelled with unique labels from $\{1, \dots, r\}$, where r is the rank of the matroid. For a basis exchange in this labelled setting, the new element gets the same label as the replaced element. More precisely, if l is the labelling function on B and $e \in B$ is replaced with $e' \in E \setminus B$, then e' gets assigned the label $l(e)$. In standard matroid terminology, a “labelled” basis is an “ordered” basis.

We consider two questions:

1. When can one labelled basis of a matroid be reconfigured to another labelled basis of the matroid using a sequence of basis exchanges?
2. What is the number of basis exchanges needed in the worst case?

Note that such labelled reconfiguration is not always possible—for example, a matroid may have a single basis B , in which case there is no way to reconfigure from one labelling of B to a different labelling.

Results. In this paper we prove that one labelled basis can be reconfigured to another if and only if for every label, the element with that label in the first basis and the element with that label in the second basis lie in the same connected component of the matroid. Equivalently, for a given basis, the permutations of labels achievable by sequences of basis exchange steps is a product of symmetric groups, one for each connected component of the matroid.

Furthermore, we prove that if one labelled basis can be reconfigured to another then $O(r^{1.5})$ basis exchange steps always suffice, where r is the rank of the matroid.

In the special case of graphic matroids, our problem is the following: given two spanning trees of an n -vertex graph, with the edges of the spanning trees labelled with the labels $\{1, 2, \dots, n-1\}$, reconfigure the first labelled spanning tree to the second via basis exchange steps. Reconfiguration is possible if and only if for every label, the edge with that label in the first spanning tree and the edge with that label in the second spanning tree lie in the same 2-connected component of the graph. This was proved by Hernando et al. [3]. In this case we can prove a tighter bound: $O(n \log n)$ basis exchange steps are always sufficient.

The rest of the paper is organized as follows. Subsections 1.1 and 1.2 contain notation and background. In Section 2 we prove the results for graphic matroids, and in Section 3 we prove the results for general matroids.

1.1 Preliminaries

For a matroid M of rank r , a *labelled basis* (or “ordered basis”) is a tuple $\mathcal{T} = (B, l)$ where B is a basis and $l : B \rightarrow \{1, \dots, r\}$ is a function that assigns a unique label to each element of B . For label i , the element with that label is given by $l^{-1}(i)$. If a basis exchange step replaces $e \in B$ with $e' \in E \setminus B$, then e' is assigned the label $l(e)$. Two bases are said to be the same if they have the same elements and the elements are assigned the same labels.

Matroid Properties. For basic definitions and properties of matroids, see Oxley [8]. In the remainder of this section, we summarize the properties that we will use. Note that we are now referring to standard unlabelled matroids.

Recall the basis exchange property of matroids:

Property 1. *For any two bases B_1 and B_2 of a matroid M and an element $x \in B_1 - B_2$, there exists $y \in B_2 - B_1$ such that $(B_1 - \{x\}) \cup \{y\}$ is also a basis.*

By Property 1 any basis of a matroid can be reconfigured into any other using at most r basis exchange steps.

The pairs of elements that can participate in exchanges can be characterized exactly. Let B be a basis. Adding an element $e \notin B$ to B creates a unique circuit C called the *fundamental circuit* of e with respect to B .

Property 2. [8, Exercise 5, Section 1.2] *Given a basis B , and an element $e \notin B$, let C be the fundamental circuit of e with respect to B . Then the set of elements e' such that $(B - \{e'\}) \cup \{e\}$ is a basis are precisely the elements $e' \in C$.*

The *fundamental graph* of basis B , denoted S_B , is the bipartite graph with a vertex for each element of the matroid, and an edge (e, f) for every $e \in B$ and $f \in E - B$ such that e, f form a basis exchange, i.e. $(B - \{e\}) \cup \{f\}$ is a basis. Note that the closed neighbourhood of $f \in E - B$ is f 's fundamental circuit.

For a matroid M with element set E and a set $T \subseteq E$, the matroid that results from contracting T is denoted M/T .

Property 3. *For matroid M , and sets $T \subseteq E$ and $X \subseteq E - T$*

$$r_{M/T}(X) = r_M(X \cup T) - r_M(T).$$

This property implies that we can augment an independent set as follows:

Property 4. *If T is independent in matroid M and A is independent in M/T then $T \cup A$ is independent in M .*

The concept of graph connectivity generalizes to matroids. For any matroid M , define the relation ψ on the elements of M by $e \psi f$ if and only if either $e = f$ or there exists a circuit of M that contains both e and f . It can be shown that the relation ψ is an equivalence relation and we say that each equivalence class defines a *connected component*. In the case of a graphic matroid, the equivalence classes are the *2-connected components* or *blocks* of the graph.

We will use Menger's theorem for graphs, see [10, chapter 9], which says that the maximum number of vertex-disjoint paths from vertex s to vertex t is equal to the minimum number of vertices whose removal disconnects s and t . We will also use the generalization of Menger's theorem to matroids, which is known as Tutte's linking theorem [11], see [8]. The statement of this theorem will be given where we use it in Section 3.2.

1.2 Relationship to Triangulations

Our present study of reconfiguring labelled matroid bases was prompted by related results on reconfiguring labelled triangulations [2].

The set of triangulations of a point set (or a simple polygon) in the plane has some matroid-like properties. In particular, given a point set P , let E be the set of all line segments d such that the endpoints of d are in P and no other point of P lies on d , and let \mathcal{I} be the set of all subsets of pairwise non-crossing segments of E . Then the set system (E, \mathcal{I}) is an independence system, but fails the third matroid axiom: if I_1 and I_2 are in \mathcal{I} and

$|I_2| > |I_1|$, there does not always exist an element e of $I_2 - I_1$ such that $I_1 \cup \{e\} \in \mathcal{I}$. (As a simple example consider five points 1, 2, 3, 4, 5 in convex position where I_1 contains segments (1, 3) and (1, 4) and I_2 contains segment (2, 5).) However, the maximal sets in \mathcal{I} all have the same cardinality—this is because maximal sets of non-crossing segments correspond to triangulations and all triangulations contain the same number of edges.

The failure of the third matroid axiom for triangulations means that the greedy algorithm does not in general find a minimum weight triangulation of a point set, and in fact the problem is NP-hard even for Euclidean weights [7].

On the other hand, triangulations share some of the reconfiguration properties of matroids—any triangulation can be reconfigured to any other triangulation via a sequence of elementary exchange steps called “flips” where one segment is removed and a new segment (the opposite chord of the resulting quadrilateral, if it is convex) is added [1]. Furthermore, there is evidence that results analogous to the ones we prove here for reconfiguring labelled matroid bases carry over to labelled triangulations [2]. One intriguing possibility is a more general result that encompasses both the case of matroids and of triangulations.

2 Reconfiguring ordered bases of a graphic matroid

In this section we concentrate on graphic matroids. We characterize when one labelled spanning tree of an n -vertex graph can be reconfigured to another labelled spanning tree of the graph. Our first result provides a bound of $O(n^2)$ exchange steps for the reconfiguration. In the second subsection we improve this to $O(n \log n)$. The first result (with the $O(n^2)$ bound) extends immediately to general matroids, but we give the details for the graphic matroid case for the sake of readers who may wish to learn only about labelled reconfiguration of spanning trees.

2.1 When are two ordered bases reconfigurable?

Theorem 1. *Given two labelled spanning trees \mathcal{T}_1 and \mathcal{T}_2 of an n -vertex graph G , we can reconfigure one to the other if and only if for each label i , the edge with that label in \mathcal{T}_1 and the edge with that label in \mathcal{T}_2 lie in the same 2-connected component of G . Moreover, when reconfiguration is possible, it can be accomplished with $O(n^2)$ basis exchange steps.*

Proof. The ‘only if’ direction is clear because an exchange of edge e to e' can be performed only if both e and e' lie in the same 2-connected component. For the ‘if’ direction, we will provide an explicit exchange sequence to reconfigure \mathcal{T}_1 to \mathcal{T}_2 .

First, pick any (unlabelled) spanning tree B and reconfigure both \mathcal{T}_1 and \mathcal{T}_2 to B while ignoring the labels. This takes $O(n)$ exchange steps. Let σ_1 be the sequence of exchanges that reconfigures \mathcal{T}_1 to B and σ_2 be the sequence that reconfigures \mathcal{T}_2 to B . Obviously, the labels of the edges of B obtained from the two sequences will not match in general. Below, we give an exchange sequence σ of length at most $O(n^2)$ to rearrange the labels in B . Thus performing σ_1 followed by σ followed by the reverse of σ_2 reconfigures \mathcal{T}_1 to \mathcal{T}_2 with $O(n^2)$ exchange steps.

The problem is now reduced to the following: given one spanning tree B and two labellings of it, $\mathcal{T}_1 = (B, l_1)$ and $\mathcal{T}_2 = (B, l_2)$, reconfigure \mathcal{T}_1 to \mathcal{T}_2 using $O(n^2)$ exchanges. We will do this by repeatedly swapping labels. More precisely, let i be a label, let $e_1 = l_1^{-1}(i)$ be the edge with label i in \mathcal{T}_1 and let $e_2 = l_2^{-1}(i)$ be the edge with label i in \mathcal{T}_2 , and suppose that $e_1 \neq e_2$. We will show that with $O(n)$ exchanges we can swap the labels of e_1 and e_2 in \mathcal{T}_1 while leaving all other labels unchanged. This moves label i to the correct place for \mathcal{T}_2 , and repeating over all labels solves the problem and takes $O(n^2)$ exchanges.

Since e_1 and e_2 lie in the same 2-connected component of G (by hypothesis), there must exist a cycle C of G that goes through both e_1 and e_2 . Let t be the number of edges of $C \setminus B$. Then $t \geq 1$. We will argue by induction on t that $6t - 3$ exchanges suffice to swap the labels of e_1 and e_2 . Let f be an edge of $C \setminus B$.

If f is the only edge of $C \setminus B$, then we can perform the swap directly: use Property 2 to exchange e_1 with f , e_2 with e_1 , and, finally, f with e_2 . This sequence returns us to B and swaps the labels of e_1 and e_2 while leaving all other labels unchanged. Thus the case of $t = 1$ can be solved with 3 exchanges.

More generally, f is not the only edge of $C \setminus B$. Adding f to B creates a cycle C' that must contain an edge $f' \in B \setminus C$. Using Property 2 we exchange f with f' . The result is a new spanning tree B' whose intersection with C is strictly increased, so $C \setminus B'$ is smaller. Also note that B' contains e_1 and e_2 . By induction, we can swap the labels of e_1 and e_2 in B' with at most $6(t - 1) - 3$ exchanges, leaving all other labels unchanged. After that, we exchange f' and f to return to B with original labels except that the labels of e_1 and e_2 are now swapped. The total number of exchanges is at most $6(t - 1) - 3 + 6 = 6t - 3$ \square

2.2 Tightening the bound

In this section we show that the $O(n^2)$ bound on the number of exchanges from the previous section can be improved. Note that the common spanning tree B we chose in the previous section, to reconfigure both \mathcal{T}_1 and \mathcal{T}_2 to, was completely arbitrary. We could have perhaps chosen a spanning tree that made the task of swapping labels easier. That is precisely what we do in this section.

Observe that it is sufficient to consider a 2-connected graph G since for a general graph we can just repeat the argument inside each of its 2-connected components. We will construct a spanning tree B of G whose fundamental graph (as defined in Section 1.1) has diameter $O(\log n)$. This proves our result based on:

Lemma 1. *Let G be a 2-connected graph, and B be a spanning tree of G whose fundamental graph S_B has diameter d . Then for any two edges of B there is an exchange sequence of length $O(d)$ that swaps the labels of those two edges while leaving other labels unchanged.*

Proof. Let e and f be two edges of B , and suppose the shortest path between them in S_B has length $t < d$. We will prove by induction on t that we can swap the labels of e and f with $3t - 3$ exchanges. Let e_1 and e_2 be the two vertices that occur immediately after e on a shortest path from e to f in S_B . Then the cycle formed by adding e_1 to B contains both e and e_2 by Property 2, and by the same property, we can perform the following exchanges:

e_1 with e , then e_2 with e_1 , then e with e_2 . This exchange sequence returns us to B while swapping the labels of e and e_2 , and leaving other labels unchanged. In the basis case of the induction $t = 2$ and $e_2 = f$ and this completes the swap with 3 exchanges. In the general case, the distance between e_2 and f in S_B is $t - 2$. By induction, we can swap the labels of e_2 and f with $3(t - 2) - 3$ exchanges. After that we repeat the first 3 exchanges to complete the swap of the labels of e and f . All other labels are unchanged. The total number of exchanges is $6 + 3(t - 2) - 3 = 3t - 3$. \square

Thus it remains to construct a spanning tree B such that the diameter of S_B is $O(\log n)$. We will construct our spanning tree by repeatedly contracting cycles. For a 2-connected graph G with edge set $E(G)$ and for a cycle $C \subseteq E(G)$, let G/C denote the graph obtained by contracting C . Note that $E(G/C) = E(G) - E(C)$. Contracting C creates *blocks* that are the maximal subgraphs of G/C that are 2-connected. In order to get a bound on the diameter of S_B we will need the following:

Lemma 2. *Any 2-connected graph G with m edges contains a cycle C such that all blocks of G/C have at most $m/2$ edges.*

Proof. Let C be the cycle that minimizes the size of the largest block obtained upon contracting it. We claim that all blocks of G/C have at most $m/2$ edges. Suppose not. Then there exists a block H of size bigger than $m/2$. We will derive a contradiction.

Some edges of $E(H)$ are incident on vertices of C in G ; let those vertices be $\{v_1, \dots, v_k\}$ in clockwise order along C . There are two paths between v_1 and v_2 along the cycle C in G , one clockwise and one counterclockwise. Let P be the one that is counterclockwise and thus contains all vertices of $\{v_1, \dots, v_k\}$. There also exists a path P' between v_1 and v_2 that uses only the vertices of H . We define C' to be $P \cup P'$. We claim that the size of the largest block of G/C' is smaller than the size of the largest block of G/C , hence reaching a contradiction.

First, note that no block of G/C' contains an edge of H and an edge not in H . For consider edges e and e' in G/C' with e in H and e' not in H . Any path from e to e' in G must go through a vertex of C , and in particular, must go through a vertex of P , since P contains all vertices of C that have an edge of H incident to them. Because C' contains all of P , therefore e and e' are in different blocks of G/C' .

Now G/C' contains two kinds of blocks: those that contain edges of H and those that do not. Blocks of the first kind must have size at most the size of H from the argument above. In fact, they must be strictly smaller than H since C' contains at least one edge of H . Blocks of the second kind must also be smaller than H since at worst, such a block contains all edges of G that are not edges of H , and there are at most $m/2$ such edges. \square

We are now ready to construct our spanning tree B .

Lemma 3. *Given a 2-connected graph G , there exists a spanning tree B such that the diameter of S_B is $O(\log n)$.*

Proof. The algorithm for constructing B proceeds in iterations $i = 1, 2, \dots$. In iteration i we will add a set B^i to B . In the first iteration we find the cycle C of Lemma 2 such that all




Figure 1: Adding a minimal set of edges P (in dashed red) to create a 2-connection between B^{i-1} (thick edges) and C^i (thin edges). P may consist of two, one, or zero disjoint paths as shown in (a), (b), or (c), respectively.

blocks of G/C are of size at most $m/2$. Let B^1 be all edges but one of C , and contract those edges (equivalently, contract C) thus breaking the graph G into several blocks. In general, in any iteration, we start with the collection of blocks produced in the previous iteration, contract a connected set of edges inside each block, and add those edges to B^i . After each contraction we eliminate loops and parallel edges. Each iteration reduces the number of vertices of G by contracting a set of edges. The algorithm terminates once the graph G is left with just one vertex.

We now describe what happens to one of the blocks that is dealt with in iteration i . For $i \geq 2$, let H^{i-1} be a block at the beginning of iteration $i - 1$, and let B^{i-1} be the edges of H^{i-1} that we contract and add to B in the $(i - 1)^{\text{st}}$ iteration. Let b be the vertex that B^{i-1} gets contracted to, and let H^i be one of the blocks formed by the contraction. Then in iteration i , we pick the edges of H^i to be contracted and added to B^i , as follows. Let C^i be the cycle of Lemma 2 for H^i . Observe that C^i may or may not include vertex b . Let e be an edge of C^i . As in the first iteration, we will add to B^i all edges of C^i except e . However, we may need to add more edges in order to maintain connectivity of B , and to ensure that S_B has small diameter.

In H^{i-1} the sets of edges B^{i-1} and C^i are disjoint. We will use the fact that H^{i-1} is 2-connected and apply Menger's Theorem, see [10, chapter 9], to find a set of edges that connect B^{i-1} and C^i . More precisely, let s be a new vertex adjacent to all endpoints of edges in B^{i-1} and t be a new vertex adjacent to all endpoints of edge in C^i . Because H^{i-1} is 2-connected, we must remove at least 2 vertices to disconnect s and t . Then, by Menger's Theorem, there are two internally vertex-disjoint paths from s to t . Let P be a minimal set of edges of H^{i-1} that form two such paths. Note that if C^i includes vertex b , one or both of the paths will have no edges of H^{i-1} . See Figure 1.

Observe that the two paths of P go from two distinct vertices that are joined by a path in B^{i-1} to two distinct vertices that are joined by a path in $C^i - \{e\}$. Thus, a cycle, D , is formed by P together with a non-empty subset of B^{i-1} and a non-empty subset of $C^i - \{e\}$. Let f be an edge of P (if P is non-empty). By minimality of P , there is no cycle in $P - \{f\}$. Add to B^i the set $(C^i - \{e\}) \cup (P - \{f\})$. This completes the description of how we handle

one block in iteration i . We handle other blocks the same way, adding further edges to B^i for each other block. This completes the description of the algorithm.

To establish the correctness of the algorithm, first note that after each iteration $B = \cup_{j=1}^i B^j$ is connected and contains no cycle. Thus, when the algorithm terminates, B spans all vertices of G and contains no cycle. It remains to prove that the diameter of S_B is $O(\log n)$. In each iteration of the algorithm, we reduce the size of each block by at least half, and thus the algorithm terminates in at most $O(\log n)$ iterations. To complete the proof we will show that for each i , every edge in B^i has a path of length $O(1)$ in S_B to some edge in B^{i-1} .

Referring to the step of the construction described above, note that the fundamental circuit of e in B^i contains all of C^i . Thus e is joined by an edge of S_B to every edge of C^i . If P is empty, then the fundamental circuit of e also includes an edge of B^{i-1} and we are done. Otherwise, the fundamental circuit of f in B contains all of D , which includes all of P together with at least one edge of C^i and at least one edge of B^{i-1} . Thus f is joined by an edge of S_B to every edge of P , and to at least one edge of C^i and to at least one edge of B^{i-1} . Therefore, in S_B every edge of B^i is within distance 4 of some edge of B^{i-1} . \square

Lemmas 1, 2, and 3 give us the following strengthened form of Theorem 1.

Theorem 2. *If \mathcal{T}_1 and \mathcal{T}_2 are two labelled spanning trees of an n -vertex graph G and for each label i , the edge with that label in \mathcal{T}_1 and the edge with that label in \mathcal{T}_2 lie in the same 2-connected component of G then we can reconfigure \mathcal{T}_1 to \mathcal{T}_2 using $O(n \log n)$ basis exchange steps.*

3 Reconfiguring ordered bases of a general matroid

In this section we turn to general matroids. We generalize the result of the previous section that characterizes when one labelled basis of a rank r matroid can be reconfigured to another labelled basis of the matroid and prove a bound of $O(r^2)$ exchange steps for the reconfiguration. In the second subsection we improve this bound to $O(r^{1.5})$, a weaker bound than was possible for graphic matroids.

3.1 Connectivity

Our goal is to follow the proof of Theorem 1, which used edge contraction, cycles, and 2-connectivity in a graph. Observe that contraction of edges in a graph corresponds to contraction of elements in a matroid, cycles in a graph correspond to circuits in a matroid, and 2-connectivity in a graph corresponds to connectivity in a matroid (every pair of elements is contained in some circuit).

With these correspondences, it is easy to check that every step of the proof of Theorem 1 goes through for matroids and thus we get the following theorem, which we state without proof.

Theorem 3. *Given two labelled bases \mathcal{T}_1 and \mathcal{T}_2 of a rank r matroid M , we can reconfigure one to the other if and only if for each label i , the element with that label in \mathcal{T}_1 and the element with that label in \mathcal{T}_2 lie in the same block of M . Moreover, when reconfiguration is possible, it can be accomplished with $O(r^2)$ basis exchange steps.*

3.2 A tighter bound for general matroids

In order to tighten the $O(r^2)$ bound on the number of exchanges needed for labelled basis reconfiguration in general matroids, we would like to follow the approach we used for graphic matroids in Section 2.2. Lemma 1 carries over directly so it suffices to build a basis B whose fundamental graph S_B has small diameter. For graphic matroids of rank r , we achieved diameter $O(\log r)$ but for general matroids we will only achieve a weaker bound of diameter $O(\sqrt{r})$. Our starting point for graphic matroids was Lemma 2 which proved that there is a cycle whose contraction cuts the size of a block in half. For general matroids the analogous result is conjectured to be true, but we must rely on the following weaker result, attributed to Seymour, and with an explicit proof in [6, Corollary 1.4].

Lemma 4. *Let C be the biggest circuit of a connected matroid M . Then the biggest circuit of M/C is strictly smaller than C .*

Using this lemma we can follow the approach we used for graphic matroids to prove the following bound.

Theorem 4. *If \mathcal{T}_1 and \mathcal{T}_2 are two labelled bases of a rank r matroid M and for each label i , the element with that label in \mathcal{T}_1 and the element with that label in \mathcal{T}_2 lie in the same connected component of M then we can reconfigure \mathcal{T}_1 to \mathcal{T}_2 using $O(r^{1.5})$ basis exchange steps.*

Proof. Following the idea of the proof of Theorem 2, we will prove the bound by showing that every connected component of the matroid has a basis B whose fundamental graph S_B has diameter $O(\sqrt{r})$. We do this using almost exactly the algorithm of Lemma 3 with one difference: instead of picking the cycle of Lemma 2 in each block in each iteration, we will pick the biggest circuit and use Lemma 4.

As before, the algorithm for constructing B proceeds in iterations $i = 1, 2, \dots$. In iteration i we will add a set B^i to B . In the first iteration we find the the largest circuit C in M and let B^1 be all elements but one of C . We contract those elements, thus breaking the matroid into several connected components. In general, in any iteration i , we start with the collection of connected components produced in the previous iteration, contract some elements inside each component, and add those elements to B^i . After each contraction we eliminate loops and parallel elements. Each iteration reduces the number of elements of M and the algorithm terminates when no elements are left.

We now describe what happens to one of the connected components that is dealt with in iteration i . For $i \geq 2$, let H^{i-1} be a connected component at the beginning of iteration $i - 1$, and let B^{i-1} be the elements of H^{i-1} that we contract and add to B in the $(i - 1)^{\text{st}}$

iteration. Let H^i be one of the connected components formed by the contraction. Then in iteration i , we pick the elements of H^i to be contracted and added to B^i , as follows.

Let C^i be the biggest circuit of H^i and let e be an element of C^i . As before, we will put $C^i - \{e\}$ into B^i , but, as before, we may need to add elements to connect the independent set B^{i-1} with the circuit C^i . We will be working in the matroid H^{i-1} which we abbreviate as H .

Since C^i is a circuit in H^i , we have $|C^i| - 1 = r_{H^i}(C^i)$. Now, H^i is a connected component of H/B^{i-1} , so by Property 3, $r_{H^i}(C^i) = r_H(C^i \cup B^{i-1}) - r_H(B^{i-1})$. Thus $r_H(C^i \cup B^{i-1}) = |B^{i-1}| + |C^i| - 1$, which means that $C^i \cup B^{i-1}$ has co-rank 1 and contains a unique circuit.

If C^i is independent in H then $C^i \cup B^{i-1}$ has a circuit formed by the elements of C^i together with at least one element of B^{i-1} . In this case we add to B^i the set $C^i - \{e\}$. Observe that this set is independent in H and that the fundamental circuit of e contains all elements of C^{i-1} and at least one element of B^{i-1} . In the graphic case, this corresponds to Figure 1 (right). We will prove below that B^i has the properties we need.

Otherwise, C^i is not independent in H . In this case, the circuit in $C^i \cup B^{i-1}$ is C^i , and we will need to add more elements to B^i in order to connect B^{i-1} with C^i . We will use the matroid analogue of Menger's Theorem which is known as Tutte's Linking Theorem [11]. This theorem applies to two disjoint sets of elements in a matroid. In our case the matroid is H , and the disjoint sets are B^{i-1} and $C^i - \{e\}$, both of which are independent in H . To ease notation, let $A = C^i - \{e\}$.

The analogue of a separating set of vertices is $\kappa_H(B^{i-1}, A)$, defined as the minimum, over sets X that contain B^{i-1} and exclude A , of $r_H(X) + r_H(E - X) - r_H(E)$. Since H is connected, this minimum is at least 1. In notation, we have:

$$\kappa_H(B^{i-1}, A) = \min_{B^{i-1} \subseteq X \subseteq E \setminus A} r_H(X) + r_H(E - X) - r_H(E) \geq 1.$$

The analogue of vertex-disjoint paths in Menger's theorem is $\square_{H/P}(B^{i-1}, A)$, defined as the maximum over sets P , of $r_{H/P}(B^{i-1}) + r_{H/P}(A) - r_{H/P}(B^{i-1} \cup A)$.

According to the version of Tutte's Linking Theorem stated as equation (8.16) in Oxley [8], there exists a set of elements P such that

$$\square_{H/P}(B^{i-1}, A) = \kappa_H(B^{i-1}, A). \quad (1)$$

As in the proof of the generalization of Tutte's Linking Theorem due to Geelen, Gerards, and Whittle [8, proof of Theorem 8.5.7], we will choose P to be a minimal set such that equation (1) holds. As shown in their proof, such a minimal P is independent, and is *skew* to B^{i-1} and A . Two sets are *skew* if the rank of their union is the sum of their ranks. In our case, B^{i-1} and A are independent, so skewness implies that $B^{i-1} \cup P$ and $A \cup P$ are independent.

Applying Property 3 to $\square_{H/P}(B^{i-1}, A)$ yields

$$\square_{H/P}(B^{i-1}, A) = r_H(B^{i-1} \cup P) + r_H(A \cup P) - r_H(P) - r_H(B^{i-1} \cup A \cup P). \quad (2)$$

When P is independent and skew to B^{i-1} and A this becomes:

$$\square_{H/P}(B^{i-1}, A) = |B^{i-1}| + |A| + |P| - r_H(B^{i-1} \cup A \cup P). \quad (3)$$

From this, it is clear that if the value of equation (1) is greater than 1, then we can delete elements of P to obtain a minimal P with $\square_{H/P}(B^{i-1}, A) = 1$.

For the remainder of the proof we define P to be a minimal set with $\square_{H/P}(B^{i-1}, A) = 1$. From equation (3) we know that $B^{i-1} \cup A \cup P$ has co-rank 1. There must be a unique circuit D in $B^{i-1} \cup A \cup P$ and D must contain all elements of P (by minimality of P) and at least one element of B^{i-1} (since $A \cup P$ is independent) and at least one element of A (since $B^{i-1} \cup P$ is independent).

Let f be an element of P and add to B^i the set $(C^i - \{e\}) \cup (P - \{f\})$. Observe that this set is independent in H . The fundamental circuit of e contains all the elements of C^i , and the fundamental circuit of f contains all the elements of P and at least one element of B^{i-1} and at least one element of C^i .

Before proceeding with the proof we will mention why the above analysis was separated into two cases depending on whether or not C^i is independent in H . Observe that if C^i is independent in H , then the minimal set P that connects B^{i-1} and $C^i - \{e\}$ is in fact $P = \{e\}$ itself, and when we choose f to be an element of P then $f = e$. It would be rather strange in this case (though strictly speaking, correct) to say that we add $(C^i - \{e\}) \cup (P - \{f\})$ to B^i . That is why we treated the case where C^i is independent in H as a separate case.

To complete the proof of the Theorem we must show that the final B , defined as $\cup B^i$, is a basis and that the diameter of S_B is $O(\sqrt{r})$. Since we continue until everything is contracted away, it is clear that B spans the matroid. Because each B^i is independent after contracting all B^j for $j < i$, Property 4 implies that B is independent in the matroid M . Thus B is a basis.

We now analyze the diameter of S_B . We first observe that the algorithm terminates in $O(\sqrt{r})$ iterations. This is because the number of possible iterations for which there exists a block containing a circuit of size $\Omega(\sqrt{r})$ can be at most $O(\sqrt{r})$ and once the size of the biggest circuit in each block has been reduced to $O(\sqrt{r})$, there can be at most $O(\sqrt{r})$ more iterations.

To complete the proof we will show that for each i , every edge in B^i has a path of length $O(1)$ in S_B to some edge in B^{i-1} . We will refer to the step of the construction described above. As noted above, the fundamental circuit of e in B contains all of C^i . Thus e is joined by an edge of S_B to every element of C^i . If P is empty, then the fundamental circuit of e also includes an element of B^{i-1} and we are done. Otherwise, as noted above, the fundamental circuit of f in B contains all of P together with at least one element of C^i and at least one element of B^{i-1} . Thus f is joined by an edge of S_B to every element of P , and to at least one element of C^i and to at least one element of B^{i-1} . Therefore, in S_B every element of B^i is within distance 4 of some element of B^{i-1} . \square

4 Conclusion

We studied the reconfiguration of labelled bases of a rank r matroid and provided an upper bound of $O(r \log r)$ on the worst-case reconfiguration distance for graphic matroids, and a bound of $O(r^{1.5})$ for general matroids. The obvious next question is whether this is tight. The only lower bound we have so far is $\Omega(r)$.

Another natural question is to find the minimum number of basis exchange steps needed to transform one given labelled basis to another. It is an open question whether this problem is NP-hard or polynomial-time solvable.

Acknowledgements

We could not have proved these results without major help from Jim Geelen. The results first appeared in the second author's PhD thesis [9]. Theorem 1 was proved jointly with Sander Vendonschot and Prosenjit Bose.

References

- [1] P. Bose and F. Hurtado. Flips in planar graphs. *Computational Geometry: Theory and Applications* 42:60–80, 2009, doi:10.1016/j.comgeo.2008.04.001.
- [2] P. Bose, A. Lubiw, V. Pathak, and S. Verdonschot. Flipping edge-labelled triangulations. *CoRR*, 2013, arXiv:1310.1166v2.
- [3] C. Hernando, F. Hurtado, M. Mora, and E. Rivera-Campo. Grafos de árboles etiquetados y grafos de árboles geométricos etiquetados. *Proc. X Encuentros de Geometría Computacional*, pp. 13-19, 2003.
- [4] J. van den Heuvel. The complexity of change. *Surveys in Combinatorics*, vol. 409, pp. 127–160. Cambridge University Press, 2013.
- [5] T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science* 412(12):1054–1065, 2011, doi:10.1016/j.tcs.2010.12.005.
- [6] N. McMurray, T. J. Reid, B. Wei, and H. Wu. Largest circuits in matroids. *Advances in Applied Mathematics* 34(1):213–216, 2005, doi:10.1016/j.aam.2004.09.002.
- [7] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *J. ACM* 55(2):11:1–11:29, 2008, doi:10.1145/1346330.1346336.
- [8] J. Oxley. *Matroid Theory, second edition*. Oxford University Press, 2011.
- [9] V. Pathak. *Reconfiguring Triangulations*. Ph.D. thesis, University of Waterloo, 2014.

- [10] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2002.
- [11] W. T. Tutte. Menger's theorem for matroids. *J. Res. Nat. Bur. Standards Sect. B* 69:49–53, 1965.